

PATENT

METHOD AND APPARATUS FOR AUTOMATIC RATE ADAPTATION IN A DOCSIS UPSTREAM

By
Yehuda Azenkot
Selim Shlomo Rakib

Background of the Invention

DOCSIS upstream channels have programmable parameters that allow the throughput on the channel to be raised and lowered, as needed by cable system operators. These parameters are controlled for the cable modem (hereafter CM) by the cable modem termination system (hereafter CMTS). These parameters include among other things: symbol rate, modulation, error correction redundancy, whether interleaving is on or off and various parameters of the interleaving, whether Trellis encoding is on or off, whether spectrum spreading is on or off, and other things specified in the DOCSIS specifications 1.0, 1.1 and 2.0. The main parameters that are programmable and which affect the throughput are the modulation and the error correction redundancy. The upstream is a noisy environment, so forward error correction and coding gains from spread spectrum transmission is useful to combat noise while providing good upstream throughput. Lower data rates are more immune to noise and other interference, whereas higher data rates are more suitable for relatively quiet channels.

The noise level in the upstream channel tends to vary during the day. Cable system operators control the CMTS to send upstream channel descriptor messages to the cable modems (CMs) to control parameters of the upstream. This is done so that the CMTS will have 10-15 dB of noise margin to receive upstream bursts even in worst case noise situations. Generally, when noise conditions worsen, the throughput of the upstreams is reduced to provide better margin for the CMTS upstream receiver so that it can successfully receive the data.

The problem is that this scenario requires the cable system operator to constantly monitor noise conditions on the upstream channels and readjust the upstream channel parameters to alter the throughput to match the noise conditions. The alternative to constant monitoring for the cable operator is to set the parameters of the upstream channels at a throughput which is low enough to provide the desired noise margin or packet loss rate in the worst case noise conditions. When the noise level is not worst case, the throughput is set lower than it need be. This is an undesirable situation.

PATENT

Therefore, it would be desirable to have a system which automatically monitors the noise level as determined by some parameter such as the packet loss rate, and automatically adjust the throughput to lower the packet loss rate to some acceptable level. This would avoid the need for cable operators to constantly monitor noise level or packet loss rate and would maximize throughput for conditions at all times since the throughput of upstream channels would be raised automatically during low noise conditions.

Summary of the Invention

There is disclosed herein a system to adjust the bit rate to the channel noise conditions automatically, or manually in some embodiments, to reduce the need for the cable system operator to monitor noise conditions on the line. The system measures the channel condition by monitoring any one or more of the following characteristics of the channel:

- 1) the packet loss rate;
- 2) the channel signal-to-noise ratio
- 3) the bit error rate
- 4) the byte error rate
- 5) channel impulse noise rate for dominant impulse noise channels.

The packet loss rate is the preferred measure of channel noise conditions for purposes of automatic or manual rate adaptation.

There are two modes of operation for rate adaptation by the system disclosed herein: manual and automatic. In the manual mode, the operator monitors one or more of the above channel characteristics, and makes any adjustments using preprogrammed burst profiles for the upstream. In other words, when the channel gets noise, the operator forces the CMTS to give upstream bandwidth grants on channels whose burst profiles have been set for a lower throughput through the use of downstream messages broadcast to all cable modems called Upstream Channel Descriptors or UCD messages. Each UCD message describes one upstream channel and sets the programmable parameters thereof. Any cable modem (CM) which receives an upstream bandwidth grant for a particular upstream channel, must transmit at the designated time in the grant using the parameters for the upstream channel defined in the pertinent UCD message.

In automatic mode, the preprogrammed burst profiles are selected according to the line noise status and the type of interference or noise present. There are two different sets of preprogrammed burst profiles depending on the channel interference type. Channel types generally fall into two interference categories: Additive White Gaussian Noise (AWGN)

channel, and Impulse Noise channel. The two sets of preprogrammed burst profiles are arranged in each group according to the bit rate and are arranged from the lowest bit rate to the maximum possible bit rate.

Brief Description of the Drawings

5 Figure 1 is a flow diagram of the preferred embodiment for a process for deciding whether a data rate change is required using two different thresholds so that hysteresis prevents excessive data rate changes.

10 Figure 2, comprised of Figures 2A and 2B, is a flow diagram for a preferred embodiment of processing in the CMTS for calculation of the packet loss percentage for each logical channel and making the determination regarding changing the data rate and changing the data rate or sending an alert to the operator to manually change the data rate.

 Figure 3, comprised of Figures 3A and 3B, is a flowchart of the CMTS processing for automatic rate adaption using SNR.

15 Figure 4, comprised of Figures 4A and 4B, is a flowchart of the process to implement automatic rate adaptation on a logical channel using bit error rate or byte error rate as an indicator.

 Figure 5 represents a flowchart for a generic process to determine the channel type, monitor noise conditions and automatically change the bit rate.

20 Figure 6 is a flow diagram for a process like that shown in Figure 5 but for manual changes in bit rate.

 Figure 7 is a table of typical burst profile parameters for different SNR values on a DOCSIS 1.x AWGN channel.

25 Figure 8 is a block diagram showing a typical cable system with DOCSIS 1.x and/or 2.0 compatible cable modems 102 and 104 sending upstream data on multiple logical channels to a DOCSIS 1.x and/or DOCSIS 2.0 compatible Cable Modem Termination System (CMTS) via Hybrid Fiber Coax (HFC) networks 108 and 110 and a combiner 112.

 Figure 9, comprised of Figures 9A and 9B, is a flowchart of the preferred process to characterize the channel noise type, set an initial burst profile and then adjust it automatically for changing noise conditions.

Detailed Description of the Preferred and Alternative Embodiments

Automatic Rate Adaptation: The channel packet loss rate is, in the preferred embodiment, measured for each logical channel separately. A logical channel is a channel

PATENT

upon which data is sent using one set of transmission characteristics, but a single physical channel, *i.e.*, radio frequency carrier signal, may have more than one logical channel or be just a single logical channel. For example, in mixed mode DOCSIS systems, a single physical channel may be divided up into different time slices, each of which has data transmitted
5 therein using a single set of transmission characteristics and each of which is a separate logical channel. For example, a mixed mode physical channel might have one time slice dedicated to transmission of Advanced Time Division Multiple Access (ATDMA) transmissions at a first symbol rate and a second time slice dedicated to transmission of digital data using Synchronous Code Division Multiple Access (SCDMA). Each of these time slices is a
10 separate logical channel. Likewise, one time slice may be SCDMA at a first symbol rate and another time slice may be SCDMA at a second symbol rate. Each would be a separate logical channel.

Packet loss rate is a measurement calculated or determined in the upstream receiver of DOCSIS cable modem termination systems (CMTS). CMTS upstream receivers are
15 responsible for making various measurements on the upstream logical channels they receive. Bit error rate, byte error rate, packet loss rate and channel SNR are measurements many CMTS upstream receivers make. Some CMTS receivers that have impulse noise detection circuitry also measure impulse noise rate on a per channel basis for channels that are of the dominant impulse noise type. These measurements are available to the automatic/manual rate
20 adaptation process for purposes of deciding whether to alter the upstream burst profile in each particular logical channel.

Hysteresis is used to alter the throughput to avoid excessive changing of burst profiles. If the packet loss rate falls below a first threshold, the data rate is increased. If the packet loss rate rises above a second threshold which is different than the first threshold,
25 then the data rate is decreased. The thresholds should be sufficiently separated to prevent excessive switching of data rates. Each logical channel has its own pair of thresholds in the preferred embodiment to take into account the characteristics of that channel.

Figure 1 is a flowchart illustrating the algorithm of the preferred embodiment to use hysteresis to moderate the alteration of burst profiles and data rate to situations where noise
30 levels have changed significantly from the noise level which caused the last change in data rate. The process of Figure 1 is repeated for each upstream logical channel when the operator has automatic rate adaptation capability turned on. In the preferred embodiment, the automatic rate adaptation is always turned on. In alternative embodiments, the operator can

PATENT

turn off automatic rate adaptation, and, if the operator does turn off automatic rate adaptation, no automatic rate adaptation or messages indicating that a rate adaptation should be made manually by the operator will be generated. If the operator does have automatic rate adaptation turned on, then the processes described herein to calculate a figure of channel merit for each logical channel and compare the figure of merit to one or more thresholds will be carried out for each logical channel, or at least the ones for which automatic rate adaptation is turned on.

Test 10 represents the process of retrieving the packet loss percentage figure from whatever process in the CMTS that calculates this value for the upstream channel being processed (or calculating this value in embodiments where the CMTS does not make this calculation). The process of test 10 then compares the packet loss percentage to a value calculated from a threshold TH1. TH1 is a programmable constant and is part of the pair of thresholds (TH1 and TH2) which define the hysteresis being used in this embodiment. The Packet_Loss_Required value is an operator programmable value in the preferred embodiment. Packet_Loss_Required is a value which defines an acceptable packet loss percentage, and is typically set at 0.5%. Test 10 determines if the packet loss percentage is greater than or equal to the value Packet_Loss_Required times $(1 + TH1)$. If it is, then noise conditions on the upstream channel being considered have exceeded the requirement to lower the data rate. Processing then proceeds to step 12 where the data rate is lowered either automatically or manually such as by step 10 providing a notification to the operator that a downward data rate adjustment needs to be made. If the data rate is to be altered automatically, step 12 sends a message to the process in the CMTS that composes downstream UCD messages telling it to lower the data rate on a specific upstream. More details on that process will be given below. After reducing the data rate in step 12, processing vectors back by path 18 to test 10 to start the monitoring process for this logical channel again. This is done because the fact that a change just occurred indicates that noise conditions on this logical channel are worth monitoring for a little while. In this preferred embodiment, monitoring of the logical channel will stop when tests 10 and 14 collectively determine that the packet loss percentage is not greater than or equal to the upper hysteresis threshold and not less than or equal to the lower hysteresis threshold and vector processing to step 20. Step 20 represents the process of switching to the next logical channel and starting the monitoring process over by vectoring processing to the beginning of the process (and changing the TH1 and TH2 values if they are different for the

PATENT

new logical channel). In some embodiments, processing is vectored to step 28 in Figure 2A from step 20 to determine if automatic rate adaptation is turned on for the newly selected logical channel before repeating the processing of steps 32 and following. In other
5 32 and following since, in this embodiment, automatic rate adaptation will be turned on for all logical channels.

If test 10 concludes that the packet loss percentage is not greater than
Packet_Loss_Required times $(1 + TH1)$, processing vectors to test 14. Test 14 evaluates the
10 packet loss percentage to determine if it is less than or equal to Packet_Loss_Required times $(1 - TH2)$. TH2 is the other of the two hysteresis thresholds. It should be set far enough from TH1 to prevent excessive switching of data rates. What is excessive switching of data rates is something which is up to the cable operator to decide.

If test 14 concludes that the packet loss percentage has fallen far enough, it vectors
15 processing to step 16 where the data rate is increased. This can be done manually by notifying the cable operator, or automatically by sending a message to the process in the CMTS that composes UCD messages informing it that a higher data rate can be sent on this particular logical channel. After increasing the data rate in step 16, processing vectors back to test 10 via path 18 to start the monitoring process for this logical channel again.

If test 14 determines that the packet loss percentage has not decreased to less than
20 or equal to the lower hysteresis threshold, processing vectors to step 20 where the next logical channel is picked, and processing is vectored back to step 10 to start the monitoring process for the next channel.

The packet loss percentage calculation is based on the Header Check Sum (HCS)
25 field and the CRC of each received packet. The HCS field is a 16-bit CRC field that ensures the integrity of the Media Access Control (MAC) header of the packet, even in a collision environment. The HCS field coverage includes the entire MAC header, starting with the FC field and including any EHDR field that may be present. The HCS field is calculated using the CRC-CCITT polynomial $(x^{16} + x^{12} + x^5 + 1)$ as defined in the ITU-T X.25 specification/standard. The CRC field is a 32-bit CRC field which ensures the integrity of the packet data unit (PDU)
30 which contains the payload data of the packet (as defined in the Ethernet/ISO 8802-3 standard). If either the HCS or CRC field processing points to an error, the packet is considered to be an error packet and is counted toward the numerator of the packet loss percentage calculation, the denominator being the total number of packets received over a

PATENT

unit of time (typically 1000 packets but other numbers will work). The packet loss percentage calculation should be calculated over a sufficiently large number of packets to prevent excessive changing of the data rate. In other words, it is undesirable to use a small number of packets for the denominator of the packet loss percentage calculation, because
5 this tends to lead to a change in data rate when channel noise conditions change over only a short time. After a switch in data rate by switching to a new burst profile, the flawed packet counter used to calculate the packet loss percentage is reset in the preferred embodiment and restarts the count.

This packet loss percentage calculation can be performed by the CMTS receiver
10 forward error correction circuitry, or the MAC process or the process of the invention which makes the decision regarding adaptation of the data rate based upon noise conditions.

The packet loss percentage calculation is based only upon packets transmitted during non contention grants. In DOCSIS, ranging bursts are packets transmitted during contention grants as are certain other types of packets such as requests for bandwidth. These types
15 of packets are not counted toward the packet loss percentage since there is a significant possibility of errors created by collisions.

In an alternative embodiment represented by dashed lines 22 and 24, processing is vectored to step 20 after completion of either step 12 or step 16 to pick another logical channel and then return to step 28 or setp 32 in Figure 2A to being the process again for
20 another logical channel. This alternative embodiment does not monitor the current logical channel being processed any further after making a change in the data rate, although all logical channels are periodically checked with one of the processes described herein to determine if a change in data rate is indicated.

Referring to Figure 2, there is shown a flow diagram for a preferred embodiment of
25 processing in the CMTS for calculation of the packet loss percentage for each logical channel and making the determination regarding changing the data rate and changing the data rate or sending an alert to the operator to manually change the data rate. This process is carried out separately for each logical channel, so block 27 represents whatever process is used to select the particular logical channel to be used. Block 27 is usually called by a subroutine call
30 from the main loop of the CMTS processor, but it could be called by any other mechanism. After selecting the logical channel to be processed, test 28 determines if the operator has turned on the automatic rate adaptation process for this particular logical channel. This is the preferred embodiment. In alternative embodiments, automatic rate adaptation is either turned

PATENT

on or off for all logical channels, and test 28 represents making this determination. Test 28 can be performed before block 27 in these alternative embodiments. In still other alternative embodiments, test 28 can be eliminated, and the automatic rate adaptation process can be on all the time for all logical channels. If the automatic rate adaptation process is not turned on, step 30 is executed to return to the calling program.

If automatic rate adaptation is turned on, step 32 is performed to reset the flawed packet counter and reset the total number of packets received counter for this particular logical channel. These counters can be implemented in either hardware or software. Next, step 34 is performed to increment the received packet counter each times a packet with a MAC header is received which is transmitted pursuant to a grant for a non contention interval. Packets transmitted for contention grant intervals like ranging packets and bandwidth request packets are not counted because they frequently suffer collisions. In the preferred embodiment, only packets that are not transmitted during contention grants are counted, so step 34 also represents the embodiments where the received packet is discarded if it is transmitted during a contention grant such as ranging. In these embodiments, the incrementation of the received packet counter and the flawed packet counter only occurs if the packet is not discarded. Step 36 is then performed to receive the packet and determine if there is an error in the packet and increment the error count. In the preferred embodiment, this processing is carried out by incrementing the flawed packet counter each time processing of the HCS or CRC field of the received packet points to an error in the packet. Step 38 is then performed to compare the count in the total number of packets received counter to see if this counter has reached the maximum packet count limit. This limit is programmable in the preferred embodiment, but may be fixed in other embodiments. Typically it is 1000 or some other fairly large number to prevent excessive switching of data rates.

Step 40 is performed when a fixed number of packets has been received and test 38 exits on the Yes path. The function of step 40 is to calculate the packet loss percentage. The calculation of step 40 takes the number of flawed packets, as determined by the flawed packet counter count, and divides this number by the total number of packets received, as determined by the count in the received packet counter.

Next, step 26 is performed to compare the packet loss percentage to the noise threshold or to two hysteresis noise thresholds if two noise thresholds are used. In the preferred embodiment, step 26 is all the processing shown in Figure 1, and two noise

PATENT

thresholds are used to implement hysteresis to prevent excessive data rate switching. In alternative embodiments, only a single noise threshold is used, and the data rate is decreased when the packet loss percentage rises above it, and increased when the packet loss percentage falls below it, and step 26 was generically worded to cover all these
5 alternative embodiments.

Test 42 represents the determination made from the comparison of step 26 as to whether the data rate needs to be changed because of the current value for the packet loss percentage. Because the flowchart of Figure 2 also represents a single threshold embodiment, step 42 is shown separately, although, in embodiments where two thresholds
10 and hysteresis are used, step 42 is really part of the processing of step 26, as shown in Figure 1.

If test 42 concludes that a data rate change is needed, step 44 is performed to send a message to the appropriate process to change the data rate. Step 46 represents the process of either automatically changing the data rate or providing a visible or audible
15 message to the cable operator to consider changing the data rate. If automatic changing of the data rate is implemented, step 46 selects a burst profile from a table of burst profiles for the next data rate up or the next data rate down, as appropriate from the comparison of the packet loss percentage to the noise threshold(s). After the new burst profile is picked, a new UCD message is prepared and transmitted to all cable modems. This UCD message
20 broadcasts the new upstream channel descriptor parameters for the logical channel being processed and includes the new data rate descriptors such as modulation type, number of parity bits, symbol rate, etc. Processing then returns to the main loop, as symbolized by path 48 returning to step 30. If test 42 determines that no change in the data rate is needed, path 48 is taken to step 30.

If manual data rate change is to be implemented, step 46 represents the process of displaying or sounding an audible alert to the operator, or both, indicating that the operator should consider increasing the data rate or decreasing it for a particular logical channel, as appropriate.

In alternative embodiments, only a single threshold is used to make the determination.
30 In still other alternative embodiments, other measures are used to make the determination of when to switch the data rate and either two thresholds separated by a hysteresis margin or a single threshold can be used. The measures that may be used to make the data throughput determination of what burst profile to assign to an upstream channel are:

- 1) packet loss rate,
- 2) bit error rate,
- 3) byte error rate,
- 4) channel SNR for AWGN channels, ingress noise and other interference type channels,
- 5) channel impulse noise rate for channels dominated primarily by impulse noise.

Channel Signal-To-Noise Ratio

The signal-to-noise ratio (SNR) of every logical channel can be measured. This measurement is done by the CMTS upstream receiver and is based on the QAM constellation error. The QAM constellation error is determined from the received constellation point versus the hard decision value of the slicer (the actual I and Q values of the ideal constellation point which was transmitted). Any variation between the hard decision value and the received constellation point is caused by noise in the channel, so the degree of variation is indicative of the SNR of the channel. The signal power used in this calculation is the nominal expected power level. In DOCSIS systems, the CMTS controls the transmission power of every cable modem (CM) in the system so that signals transmitted from that CM arrive at the CMTS at the nominal power.

The automatic rate adaptation system can use the SNR to change the logical channel throughput also in this alternative embodiment by calculating the SNR, comparing it to some figure of merit and selecting a burst profile with the appropriate throughput.

Figure 3, comprised of Figures 3A and 3B, is a flowchart of the CMTS processing for automatic rate adaptation using SNR. As was the case for the flowchart of step 2, steps to selected a logical channel and determine whether automatic rate adaptation is turned on for the logical channel can be added to the flowchart of Figure 3 for some embodiments. Only the steps to process a single logical channel assuming automatic rate adaptation is turned on are shown. The total packets received counter and flawed packets counter are eliminated in this embodiment, and a mechanism to calculate an average signal-to-noise ratio over a time interval or window is substituted. Step 50 represents the process of resetting a signal-to-noise (SNR) averaging mechanism for this logical channel. Step 52 increments the received packet counter each time a non contention packet is received. In the preferred embodiment, only packets that are not transmitted during contention grants are counted, so step 52 also represents the embodiments where the received packet is discarded if it is transmitted during a contention grant such as ranging. In these embodiments, the incrementation of the received packet counter only occurs if the packet is not discarded. Step 54 calculates the

PATENT

running average_noise based upon a burst_noise value output by PHY chip in the CMTS receiver. Every ranging burst has a burst_noise value output by conventional PHY chips. The PHY chip of the invention will have its circuitry modified to output an average_noise measurement for every burst or at least for every non contention burst. The average_noise measurement is based upon the average variance between the I and Q values of the received constellation points that have been noise corrupted versus the ideal I and Q values of the constellation point that would have been received in a perfect channel where noise corruption did not occur (the original constellation point transmitted). The noise power in a noise corrupted constellation point is proportional to the square of the distance between the I value received and the ideal I value averaged with the square of the distance between the Q value received and the ideal Q value. The average_noise measurement is output on a per burst basis. The running Average_noise would be the running sum of average_noise values divided by the total number of average_noise values taken into account in the calculation (the total number of average-noise measurements added together by the end of the window interval). In alternative embodiments, the average_noise value is an analog signal which is passed through a low pass filter to get the average value. The level at the output of the low pass filter at the time the averaging interval ends is the average_noise for that interval.

Step 56 determines if the max received packet count has been reached signifying the end of the interval over which average SNR is to be calculated. If not, processing returns to step 52 to continue to receive packets. If the max packet count has been reached, step 58 calculates the average SNR of the logical channel over the interval defined by the maximum packet count using the average variances of the received constellation points from the hard decision thresholds that define the actual I and Q values of the transmitted constellation point. The average SNR will be the nominal signal power divided by the square of the average_noise over the interval.

Step 60 compares the average SNR of the logical channel to one or more constants calculated from a programmable SNR_REQUIRED value selected for the burst profile in use for the logical channel as modified by an SNR threshold. In the preferred embodiment, an SNR threshold TH1 is subtracted from the value SNR_REQUIRED and a SNR threshold TH2 is added to the SNR_REQUIRED value. Although hysteresis and two SNR thresholds are the preferred embodiment, in alternative embodiments, only a single threshold is used. Step 60 comprises the following algorithm in the preferred embodiment:

PATENT

(1) IF average SNR < average SNR_REQUIRED * (1-TH1), Then Reduce The Data Rate Throughput, Preferably According To A Preprogrammed Burst Profile Table By Picking The Next Data Rate Down;

5 (2) IF average SNR > average SNR_REQUIRED * (1 + TH2), Then Increase The Data Rate, Preferably According To Preprogrammed Burst Profile Table By Picking The Next Data Rate up;

(3) End.

TH1 and TH2 are threshold average SNR ratios separated by a sufficient amount to provide a hysteresis amount which is sufficient to prevent excessive data rate switching.
10 average SNR_REQUIRED is the required amount of average SNR for acceptable transmission quality and error rate for the specific burst profile in use.

In alternative embodiments, only a single SNR threshold may be used.

Step 62 determines if there is a need to change the data rate based upon the comparisons done in step 60. If not processing returns to the calling function, as
15 represented by step 70. If so, step 64 is performed to send a message to the appropriate process to change the data rate.

Step 66 then is performed to select a new burst profile as appropriate to increase the data rate throughput. In the preferred embodiment, the data rate is increased up by one step or down by one step, as the case may be in a burst profile table. The term step is intended
20 to mean one step in a burst profile table that lists burst profiles in ascending or descending order of data throughput rate. In the preferred embodiment, after the new data rate is automatically selected, an UCD message is prepared that defines the logical channel being processed and includes its new data throughput rate. This UCD message is broadcast to all CMs or at least to those which might be assigned to the logical channel. In manual systems,
25 step 66 represents generating a visible or audible alarm to the operator suggesting a data rate change. As is the case for Figure 2, step 66 represents picking a new data rate or burst profile and suggesting this data rate be adopted for a new UCD message for this logical channel.

Referring to Figure 4, comprised of Figures 4A and 4B, there is shown a flowchart of
30 the process to implement automatic rate adaptation on a logical channel using bit error rate or byte error rate of the logical channel as an indicator. The process starts at 72 and transitions to step 74 where a total codeword counter/accumulator is reset to zero. The total codeword counter/accumulator sums the total number of codewords received during the

PATENT

pertinent interval. Step 74 also resets a total flawed codeword counter/accumulator to zero. The total flawed codeword counter/accumulator sums the total number of codewords with uncorrectable errors during the pertinent interval. Each counter/accumulator is reset to zero for the particular logical channel being processed.

5 When each burst is received by the CMTS, the R-S decoder in the PHY circuitry makes three measurements for each burst which are output to the MAC process. The number of good codewords with no errors is output. The number of codewords that had errors that were corrected is output. And finally, the number of codewords with uncorrectable errors is also output. Each codeword has a plurality of bytes which depends
10 upon the codeword size.

 Step 76 represents the process of receiving a packet and adding the total number of codewords received to the running total of codewords received for this logical channel for the interval being processed. This can be done in an accumulator in hardware or by a software process to keep the running total. This is done only if the packet is not transmitted
15 in response to a contention grant. Any packets received which were transmitted in response to a contention grant are discarded in the preferred embodiment, and only packets with specific IUCs are counted. In the preferred embodiment, only packets with the IUC of a long data packet are counted, but in alternative embodiments, the total number of codewords received in both long and short data packets is counted. The running total of codewords
20 received for this logical channel at the end of the interval is the quantity n_{total} in equation (1) below.

 Step 78 represents the process of adding the number of codewords reported to have uncorrected errors to the running total of codewords with uncorrectable errors for this logical channel for the interval being processed. This quantity is the variable n_{error} in
25 equation (1) below. The reference to a specific IUC in the definition of n_{error} below is a reference to the fact that in the preferred embodiment, only the number of uncorrectable codewords in a burst having the IUC of a long data packet are counted. In other embodiments, other non contention IUCs can be included.

 Step 80 represents the process of determining if the desired interval has elapsed.
30 This can be done by comparing the total number of packets has been received to a programmable number of maximum packets to process before making a judgment as to whether the data rate of the channel should be increased or decreased. This can also be done by counting the total number of codewords or bytes received.

PATENT

In the embodiment represented by Figure 4, the judgment as to whether to increase or decrease the data rate is based upon the bit error rate or byte error rate of the logical channel. Step 82 represents the process of calculating or obtaining the bit error rate or byte error rate of the logical channel being processed. In the preferred embodiment, the MAC process running in the CMTS determines the bit error rate or byte error rate by evaluating equation (1) below. In other embodiments, the rate adaptation process calculates the bit error rate or byte error rate itself by evaluating equation(1) below.

$$(1) \text{ ByteErrorRate} = \frac{(T + 1) * n_error}{(k + 2T) * n_total}$$

where

T: the number of maximum correctable bytes in an Reed-Solomon (hereafter R-S) codeword.
T+1: when there is an uncorrectable R-S codeword, there is a high probability of having T+1 flawed bytes.

n_error: the total number of uncorrectable R-S codewords over the pertinent interval. This is counted for a specific IUC burst type over an interval determined by step 80 in the preferred embodiment. The number of uncorrectable codewords per burst are output as a measurement made by the R-S decoder in the PHY chip and is sent to the MAC process where a running sum over the pertinent interval is kept.

n_total: the total number of received R-S codewords within a certain time interval. In the preferred embodiment, on the total number of codewords for a specific IUC burst are counted, but in alternative embodiments, the total number of codewords for multiple IUC burst types of a non contention nature are counted.

k: the number of information bytes in an R-S codeword of a specific IUC; and

k+2t: the R-S codeword length in bytes.

Since on average, only half of the bits are erroneous, the bit error rate BER is given by:

$$(2) \text{ BER} = \frac{0.5 * (T + 1) * n_error}{(k + 2T) * n_total}$$

Generally, but not in all embodiments, the long data grant IUC is used the calculate the BER or the byte error rate. Step 82 represents the process of calculating the BER or byte error rate by evaluating the appropriate one of formulas (1) or (2) either in the CMTS MAC process or the automatic rate adaptation process itself.

Step 84 represents the process of comparing the bit error rate or byte error rate calculated in step 82 to one or more thresholds. Assuming a single threshold is used, if the BER or byte error rate is above the threshold, the data rate will have to be reduced. If the BER or byte error rate is below the threshold, the data rate will have to be increased.

Step 86 makes the determination as to whether the data rate must be changed based upon the results of the comparison. If a conclusion is drawn that the data rate must be changed, step 88 is performed to send a message to the appropriate process to change the data rate. This step may be eliminated in some embodiments, as is true for the identical step in Figures 2B and 3B, and processing may proceed directly to actually changing the data rate, as symbolized by step 90. Step 90, as is the case for the identical steps in Figures 2B and 3B, represents any way of changing the data rate in the appropriate direction based upon the comparison to the threshold(s). In the preferred embodiment, this is done by selecting the next burst profile up or down, as appropriate, from the current burst profile in a table that lists burst profiles in ascending or descending order of data rates. Step 90 also represents the process of generating and sending a new UCD message for the logical channel being processed to set new transmission parameters for the logical channel which includes parameters to implement the new data rate. If step 86 determines that no new data rate is necessary, processing flows to step 92 to return to the main loop of processing. Typically, the main loop will select a new logical channel, and call the process again by vectoring processing to step 72.

INITIAL PROCESSING TO DETERMINE CHANNEL TYPE AND SELECT INITIAL BURST PROFILE

Some channels have a large amount of impulse noise. Other channels are primarily average white gaussian noise (AWGN). The burst profile that should be selected for each different type of channel differs based upon whether the channel is AWGN or has more impulse noise. The optimum burst profile for an AWGN channel is not suitable for an impulse noise channel because reducing the modulation level to improve the BER for AWGN might do the opposite for an impulse noise channel when the R-S code is weaker. Therefore, in the preferred embodiment, two sets of burst profiles are used, each set optimized for a different type of noise on the channel. One set is optimized for AWGN channels, and the other is optimized for impulse noise channels. Each set is included in a table and sorted by ascending order of bit rate in steps. The burst parameters of modulation type, Error Correction Code parameters (k, T for R-S error correction codes, and Trellis coding enabled

PATENT

or disabled for Trellis Code Modulated (TCM) bursts, as well as the preamble length and the number of active codes for Synchronous Code Division Modulated bursts define the bit rate of a burst profile.

For channels with a large amount of impulse noise, the R-S code rate should be kept low. The symbol rate and carrier frequency burst parameters may also be adjusted in impulse noise channels to improve the performance in the following cases:

- 1) impulse noise present and a low SNR (in such a case, a lower symbol rate is more robust in the presence of impulse noise interference and AWGN);
- 2) narrow band interference or ingress noise (in such a case, a change in the symbol rate or the carrier frequency can shift the spectrum sufficiently out of the interference to improve the performance).

The burst profiles are selected automatically in some embodiments, and manually in others, but in each case, the burst profile should be selected according to the type of channel noise which is prevalent. The generic process to do this is described next.

Figure 5 represents a flowchart for a generic process to determine the channel type, monitor noise conditions and automatically change the bit rate. In automatic mode, the process of burst profile selection starts first with a determination of whether the channel is AWGN or impulse noise, as symbolized by step 84. This is one of the first things that happens in the CMTS boot process after powerup step 82. A channel may be characterized as AWGN or impulse noise by any process that works. The channel SNR and the impulse noise length and rate are good tools to estimate the line noise conditions. The channel SNR can be estimated by measuring the received input signal power during upstream minislots that not have been awarded to any cable modem (meaning no CM is transmitting during those minislots). The SNR can then be calculated by dividing the nominal power the CMTS commands CMs to use when transmitting by the measured noise power represented by the input power received when no CM was transmitting.

The existence of impulse noise on the channel can be determined in step 82 by monitoring the erasure output signals from the impulse noise detector. In the CMTS, an impulse noise detector circuit detects the presence of impulse noise and outputs erasure indications for samples of input signals from a logical channel that may be affected by impulse noise. The erasure percentage can be used to determine whether a channel is characterized by impulse noise. The impulse noise detector in the CMTS receiver outputs an erasure signal which is kept active for the duration of the detected impulse noise. Impulse

PATENT

noise can be detected in accordance with the teachings of US patent application 10/255,286, filed 9/25/2002, which is hereby incorporated by reference. If the number of erasures in a window of time is smaller than a predetermined threshold, the channel is determined to be an AWGN channel.

5 The impulse noise detector in the CMTS has three counters called "TDMA_erasure_count", "SCDMA_erasure_count" and "sub_symbol_count". All three counters are 32 bits in width. The "TDMA_erasure_count" counter counts the number of erasure samples in the TDMA channel over a specific interval, which means the number of samples that have an erasure bit set. The "SCDMA_erasure_count" counter counts the number of erasure samples in the SCDMA channel over a specific interval. The "sub_symbol_count" counter counts the number of sub_symbols over the same specific interval. The erasure percentage for TDMA and SCDMA logical channels is defined by the following equations.

10 (3)
$$\text{TDMA_erasure_percentage} = \frac{\text{TDMA_erasure_count}}{\text{sub_symbol_count}} * 100$$

 (4)
$$\text{SCDMA_erasure_percentage} = \frac{\text{SCDMA_erasure_count}}{\text{sub_symbol_count}} * 100$$

 In the preferred embodiment, the determination as to whether a logical channel is dominated by impulse noise is determined by the following monitoring algorithm to switch over to a characterization of a channel as an impulse noise dominated channel if the logical channel is currently deemed an AWGN channel:

20 if TDMA_erasure_percentage > TDMA_erasure_percentage_required*(1+TT1)
 then TDMA channel type = Impulse Noise
End

 TT1 is a threshold used to establish the upper limit of TDMA_erasure_percentage which is tolerated before a logical channel is declared to be an impulse noise dominated channel.

 In the preferred embodiment, the determination as to whether a logical channel is dominated by AWGN is determined by the following monitoring algorithm to switch over to a characterization of a channel as an AWGN dominated channel if the logical channel is currently deemed an impulse noise channel:

30 if TDMA_erasure_percentage < TDMA_erasure_percentage_required*(1-TT2)

then TDMA channel type = AWGN

End

TT2 is a threshold used to establish the lower limit of TDMA_erasure_percentage which is tolerated before a logical channel is declared to be an AWGN dominated channel.

5 TT1 and TT2 are separated by enough distance to provide a hysteresis effect to prevent excessive switching between channel characterization. Default values for TT1 and TT2 are TT1 = 0.3 and TT2 = 0.3. TDMA_erasure_percentage_required is a programmable constant in the preferred embodiment, but a fixed constant in other embodiments. A typical default value is 5.

10 For SCDMA channels, the determination of dominant noise type on a logical channel is determined in a similar fashion using the following monitoring equations and algorithm. If the SCDMA channel is deemed AWGN, monitoring to recharacterize the channel as impulse noise dominated is done by evaluating the expression and the following algorithm

15 if SCDMA_erasure_percentage > SCDMA_erasure_percentage_required*(1+ST1)
then SCDMA channel type = impulse noise

End

If the SCDMA channel is deemed impulse noise, monitoring to recharacterized the channe as AWGN is done by evaluating the expression and the following algorithm

20 if SCDMA_erasure_percentage < SCDMA_erasure_percentage_required*(1-ST2)
then SCDMA channel type = AWGN

End

ST1 and ST2 are constants used for hysteresis to prevent excessive switching of channel characterization. Default values are ST1=0.1 and ST2 equals 0.1, and

25 SCDMA_erasure_percentage_required is a programmable constant in the preferred embodiment, but a fixed constant in other embodiments.

Once the dominant noise characteristic of the channel is determined, the set of burst profiles for AWGN or impulse noise channels is selected, as appropriate, as symbolized by step 84.

30 Next, as symbolized by step 86, an initial burst profile is selected from the set of burst profiles adapted for the particular type of logical channel noise conditions present. The burst profile can be any profile, but preferably will be the slowest bit rate in some embodiments where monitoring of noise conditions on each channel occurs frequently and periodically so that the bit rate can be raised soon. In the preferred embodiment, the initial

burst profile is selected based upon the calculated SNR from step 84 with high SNRs causing a high bit rate to be selected initially and low SNRs causing a low bit rate to be selected initially.

Then, as symbolized by steps 88 and 90 and 92, noise conditions are continually
5 monitored, and the appropriate adjustments to the burst profile are made by selecting from the appropriate set of burst profiles according to the packet loss percentage, the SNR, the BER or the byte error rate. The processes to calculate packet loss percentage, SNR, BER or byte error rate and compare the calculated result to thresholds for making a decision
10 regarding changing the bit rate have been previously described. Any one of these processes can be used for step 88 to calculate the quality of channel parameter indicative of current noise conditions. Step 90 represents the process of comparing the quality of channel parameter so calculated to one or more thresholds, and step 92 represents the process of evaluating the comparison results and drawing a conclusion regarding whether a change in bit rate is recommended.

15 If the packet loss percentage, BER or byte error rate is too high, or the SNR is too low, the next burst profile down in the table of burst profiles selected in step 84 is selected, as symbolized by step 94. In other words, the next burst profile down with a lower bit rate will be selected by step 94 if noise conditions have increased. If the packet loss percentage, BER or byte error rate is lower than the pertinent threshold, or the SNR is higher than the
20 threshold, the next burst profile up in the selected table with a higher bit rate is selected in step 94 to take advantage of improved noise conditions. Step 94 also automatically generates and sends a UCD message so that all CMs using the logical channel for which the bit rate was just changed will start using the new burst profile.

After step 94, step 96 returns processing to whatever other processing in the boot
25 process is required. If step 92 concludes no change in bit rate is needed, step 96 is performed also. In some embodiments where multiple upstream logical channels are in use, step 98 is optionally performed after or before step 96 to select a new logical channel and repeat the process for this new logical channel.

Figure 6 is a flow diagram for a process like that shown in Figure 5 but for manual
30 changes in bit rate. The process of Figure 6 automatically determines the dominant noise type on the channel and picks the appropriate set of burst profiles for the dominant noise type, and automatically monitors noise conditions, and when a data rate change is indicated, generates a message indicating to an operator that a data rate change is advisable. The

PATENT

difference between the process of Figure 5 and Figure 6 is embodied in step 100 where instead of automatically picking a new burst profile and generating a UCD message, a message is generated to the cable operator suggesting a bit rate change.

The optimal burst profiles for impulse noise channels generally require shorter R-S
5 codewords. The optimal burst profiles for impulse noise channels have the following characteristics:

(1) R-S information symbol k is small and is adjusted according to the number of erasures in an RS codeword.

(2) The R-S interleaver depth, used only for advanced TDMA logical channels in DOCSIS 2.0
10 compatible systems, is large (according to the impulse noise duration), and the RS interleaver length is $k + 2t$. In dynamic mode, the R-S interleaver size B_r has to be set accordingly.

(3) The number of spreading intervals per SCDMA frame in DOCSIS 2.0 compatible systems is large.

Figure 7 is a table of typical burst profile parameters for different SNR values on a
15 DOCSIS 1.x AWGN channel.

Figure 8 is a block diagram showing a typical cable system with DOCSIS 1.x and/or
2.0 compatible cable modems 102 and 104 sending upstream data on multiple logical channels to a DOCSIS 1.x and/or DOCSIS 2.0 compatible Cable Modem Termination System (CMTS) via
Hybrid Fiber Coax (HFC) networks 108 and 110 and a combiner 112. A physical layer
20 demodulator 114 provides upstream data and measurements to a media access control process executed by MAC processor 116. The impulse noise detector circuitry is within demodulator 114. The processes described herein execute as part of the MAC process run by MAC processor 116. The downstream path is not shown, but the UCD messages are sent to the CMs on a downstream path and downstream logical channels which are at least
25 frequency division multiplexed from the upstream logical channels.

Currently offered CMTS from the assignee of the present invention offer 44 different pre-designed modulation profiles. They are categorized as follows: 9 for DOCSIS 1.x TDMA mode; 13 for DOCSIS 2.0 Advanced Time Division Multiple Access (ATDMA) mode; 13 for DOCSIS 2.0 Synchronous Code Division Multiple Access (SCDMA) mode; and 9 for ATDMA-
30 TDMA Shared mode. According to the teachings of the invention, these burst profiles are sorted from maximum noise robustness (minimum throughput capacity) to minimum noise robustness (maximum throughput capacity).

PATENT

Figure 9, comprised of Figures 9A and 9B, is a flowchart of the preferred process to characterize the channel noise type, set an initial burst profile and then adjust it automatically for changing noise conditions. Step 118 represents powering up the MAC process in the CMTS and measuring the SNR of a logical channel before sending any downstream grants giving any CM permission to transmit on that channel. In the case of mixed mode channels (such as SCDMA on one logical channel during first intervals and TDMA on another logical channel in different intervals), the SNR of each logical channel is measured.

Step 120 represents selecting an initial burst profile from the appropriate set of burst profiles for the channel based upon the mode for that channel (TDMA, ATDMA, SCDMD or ATDMA-TDMA Shared) and the computed SNR. In the case of mixed mode channels, an initial burst profile is selected from the appropriate set for each logical channel.

Step 122 represents monitoring the channel for its SNR and the rate at which the impulse noise detector detects impulse noise. Step 124 represents the process of comparing the impulse noise erasure rate to a threshold to determine if the channel should be characterized as an impulse noise channel or an AWGN channel. If the erasure rate is above the threshold, the channel is declared an impulse noise dominated channel in step 126, and an appropriate set of burst profiles is selected. If the erasure rate is below the threshold, step 128 declares the channel an AWGN dominated channel and selects an appropriate set of burst profiles.

Step 130 represents the process of determining the packet error rate by monitoring the HCS and CRC frame error statistics developed by the error detection circuitry in the MAC processor 116.

Step 132 represents the process performed after either step 126 or 128 is performed of comparing the packet error rate to a programmable packet error rate (PER) threshold previously entered by the cable operator, typically 0.1%. In alternative embodiments, the packet error rate threshold is a fixed number.

If step 132 determines the packet error rate is greater than the PER threshold plus a hysteresis factor, the CMTS automatically decides to increase the noise robustness by selecting the next lower bit rate burst profile in the set appropriate to the dominant noise type, as symbolized by step 134. If the channel is AWGN, the CMTS is free to select any other burst profile in the set of burst profiles including burst profiles specifically designed for impulse noise channels. If however, the channel is dominated by impulse noise, the CMTS must pick the new burst profile with a lower bit rate from only the subset of burst profiles

PATENT

that have been designed to provide noise immunity in the presence of impulse noise for the mode of the channel. If step 132 determines that the PER is less than the PER threshold minus a second hysteresis factor, the CMTS will decide to decrease the noise robustness by selecting the next higher bit rate burst profile in the appropriate set. If the channel is AWGN, the burst profile may be selected from any burst profile in the set, but if the channel is dominated by impulse noise, the new burst profile must be selected from the set of burst profiles designed to give noise immunity in the face of impulse noise for the mode of the channel.

The hysteresis factors insure that the changes in PER must be significant before a change in the burst profile will occur.

It is important that the changes in bit rates cure the PER problem and bring the PER within the desired window of acceptable PER rates. Steps 136 through 140 represent the process of monitoring the effect of repeated attempts to cure the PER problem, and jump to a new channel if the PER rate is not brought within limits. Step 136 represents the process of repeating the PER monitoring process starting with step 122 after the bit rate has been changed and continuing to perform the process and change the bit rate until either the PER is brought within limits or the maximum number of attempts has been made. Step 138 represents the process of determining if the maximum number of attempts to cure the PER problem have been made. If not, step 136 is repeated. If a pass through the process results in success in bringing the PER within limits, path 137 is taken to step 139 to resume other boot processing. Alternatively, step 141 may be performed to pick another logical channel and repeat the process of steps 118 (to calculate the SNR only of the new logical channel) through 140 for another logical channel. If the maximum number of attempts has been made, step 140 is performed to switch to a new logical channel by sending new MAP messages to the CMs that have been using the logical channel in response to bandwidth requests therefrom and sending a new UCD message if necessary to define the logical channel and desired burst profile the CMs are being instructed to switch to in the MAP messages.

Although the invention has been disclosed in terms of the preferred and alternative embodiments disclosed herein, those skilled in the art will appreciate possible alternative embodiments and other modifications to the teachings disclosed herein which do not depart from the spirit and scope of the invention. All such alternative embodiments and other modifications are intended to be included within the scope of the claims appended hereto.